

Where to look first? Behaviour control for fetch-and-carry missions of service robots

Markus Bajones¹, Daniel Wolf¹, Johann Prankl¹, Markus Vincze¹

Abstract—This paper presents the behaviour control of a service robot for intelligent object search in a domestic environment. A major challenge in service robotics is to enable fetch-and-carry missions that are satisfying for the user in terms of efficiency and human-oriented perception. The proposed behaviour controller provides an informed intelligent search based on a semantic segmentation framework for indoor scenes and integrates it with object recognition and grasping. Instead of manually annotating search positions in the environment, the framework automatically suggests likely locations to search for an object based on contextual information, e.g. next to tables and shelves. In a preliminary set of experiments we demonstrate that this behaviour control is as efficient as using manually annotated locations. Moreover, we argue that our approach will reduce the intensity of labour associated with programming fetch-and-carry tasks for service robots and that it will be perceived as more human-oriented.

I. INTRODUCTION

The ability of a robot autonomously interacting with a human requires sophisticated cognitive skills including perception, navigation, decision making and learning. Impressive achievements have already been made in the research field of HRI considering robots as tour guides in museums [1], [2], shopping malls [3] and also for assistive robots in the care and domestic context [4].

However, one of the biggest challenges still is the integration of methods into operational autonomous systems for the domestic context, which achieve satisfying results for their end-users. Often a miss-match between user expectations and robot performance can be observed in HRI studies [5], for instance because the robot behaviour is not legible to the users or is simply too slow. Imagine you command your household robot to bring you your mug and the robot could only search for it at pre-programmed places: What if the mug is at none of these places or the user rearranged the flat? Another solution could be that the robot would navigate to the center of the room, rotate several times and then start a time-consuming brute-force object detection everywhere in the map. This behaviour would already increase flexibility, but would still not be very legible for the user: Why getting an overview and then start a time-consuming search?

Would it not be much more intelligent and legible if the robot first gets an overview of the environment and then looked at the most probable locations for the mug to be, e.g. on tables or in the cupboard? One way to gain the information about this kind of relations is to extract it from knowledge databases.

Our solution for this problem is to develop a flexible behaviour controller for fetch-and-carry tasks implemented with SMACH, a Python-based library for building hierarchical concurrent state machines, as pictured for our proposed implementation in Fig. 1. We present a framework that is able to generate locations on-the-fly without the need of pre-learned object-location relationship and show in a first experiment in an ambient-assistive-living lab to proof the robustness of our approach.

The remainder of this paper is structured as follows: First, we review related research on service robots performing fetch-and-carry tasks in Section II. We then outline how the SMACH framework works in Section III. Each of the modules and their integration in the behaviour controller is then described in Section IV with a focus on the modules, but not the individual methods. Finally, a first experiment in a living room setting is presented in Section V. A summary of lessons learned from the experiment and an outlook on future work conclude the paper.

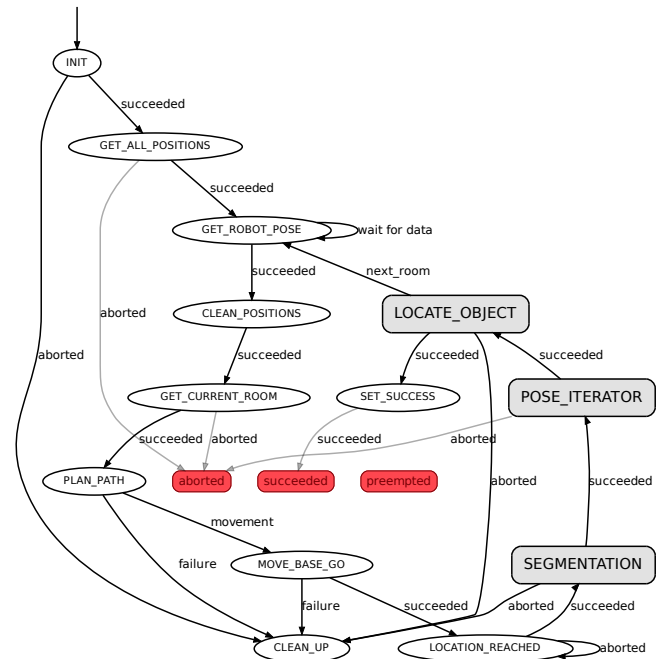


Fig. 1: Simplified state machine representation of our “fetch-and-carry” object task approach. Grey states denote sub-state machines.

¹Faculty of Electrical Engineering and Information Technology, Vienna University of Technology, 1040 Vienna, Austria {bajones, wolf, prankl, vincze}@acin.tuwien.ac.at

II. RELATED WORK

Object search and fetch-and-carry tasks represent one of the most common tasks of service robots. In traditional fetch-and-carry tasks the robot "simply" has to pick up an object from a known location and deliver it to the user. For instance, in the RoboCup@Home competition 2008/9 [6] this has been one of the major tasks. The complexity of this task can vary, e.g. the robot has to identify the correct object from a set of objects or in a cluttered scene or the robot gets a very concrete instruction from where to fetch it, such as "Bring me the mug from the kitchen table". Thus, fetch-and-carry also includes object recognition and autonomous grasping, popular examples are e.g. [7] and [8].

Recently, the perception part has come more into the focus of fetch-and-carry tasks by adding the search for object components. A popular example is George, the curious attentive semantic robot [9]. The idea of George is similar to that presented in this paper: Make a robot's search behaviour more intelligent. An intelligent search procedure is more efficient and therefore the fetch-and-carry task become more satisfying for the end-user. [10] and [11] have already shown that there is the need to make robot's behaviour more efficient and legible for the user.

Therefore, the aim of the work presented in the following is to enable a domestic service robot to perform fetch-and-carry tasks in a user-satisfying manner. To that end, a cognitive robotic system needs to integrate multiple different subsystems. Different tools, libraries and frameworks to develop robotic architectures already exist, such as RobletRTechnology by Baier et al. [12], graphical tools such as Choreographe and NAOqi for Nao robots [13], the behaviour Markup Language (BML) [14] and the Robot Operating system ROS [15]. A good overview and discussion on the advantages and disadvantages of all these approaches is given in [16].

In the following we will present SMACH, a Python-based framework that allow ROS module integration, we used to develop a behaviour control combining intelligent object search, grasping objects, and bringing them to the user.

III. SMACH

SMACH is a Python-based library for building hierarchical concurrent state machines, which also provides a ROS integrated module to design and execute simple tasks as well as complex robot behaviours. SMACH provides the possibility to quickly create prototype state machines by reusing Python design patterns. Within SMACH a task is defined by the following elements:

- **State** represents a blocking execution with pre-defined outcomes. The result of a state specifies the transitions to the next states.
- **Container** is a set of one or multiple states and defines their final outcome based on the integrated states. Especially important and useful are the StateMachine and the Concurrence containers. In the StateMachine container all the states are executed one at a time, whereas in a Concurrence multiple states are executed in parallel.

- **Transitions** define, given an outcome, to which state or container the execution chain should turn to.
- **Userdata** is data which can be shared between multiple states and containers, allowing to adapt the state outcome based on accumulated data from previously executed states.

For seamless integration SMACH provides interfaces to the three communication methods available in ROS, i.e. messages, services and Actionlib. Provided is the MonitorState for listening to published topics, executing a service represented as a state with ServiceState, calling an Actionlib interface within SimpleActionState as well as the possibility of wrapping a state machine inside an Actionlib server. With SMACH viewer, a tool to debug and visualize the running state machine, the provided user data as well as the currently executed states and containers, is also integrated.

IV. IMPLEMENTATION

For the fetch-and-carry scenario we compare two different implementations within a ROS environment and the SMACH state machine architecture. Both methods rely on a SLAM based 2D map of our laboratory in which we annotate virtual rooms as seen in Fig. 2 to create a multiple room scenario.

A. Manual definition of search positions

Starting from the SLAM-based map we manually place "search positions" at user-defined locations inside the map. To determine the pose at which the robot begins its task we define a cost function

$$c(x, y) = c_{bat}(y) - k_1 c_{prob}(x, y) + k_2 c_{pen}(y) \quad (1)$$

modelling the cost to search for a given object x at each possible search location y . $c_{bat}(y)$ represents the cost of the battery usage estimated by the path length to the pose associated to search location y . $c_{prob}(x, y)$ is the probability of object x being found at location y transformed to a cost and is initialized with a uniform distribution across all search locations. The constant factors k_1 and k_2 are introduced to normalize the cost levels. $c_{pen}(y)$ is a penalty term depending on the room the user is currently in:

$$c_{pen}(y) = \begin{cases} k_{pen} & y \text{ in same room as user} \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

The penalty term is based on the assumption that the user does not ask the robot to look for objects which are in the same room as the user himself. It ensures that the robot does not search in this room unless the object has not been detected at the locations in all other possible rooms. During the search task all the objects, which can be recognized, and are localized at a certain position trigger the adaptation of c_{prob} . Through this simple method Fig. 3 we will learn in course of time where the user places objects and where the robot should start its search. This provides a way to learn the location of specific objects based on the users preference and does not leverage upon a pre-defined knowledge database or data mining from online sources.




Fig. 2: Virtual room arrangement in our ambient-assistive-living lab.

```

while object located is FALSE do
  for all search positions do
    Calculate path length  $c_{obj}$ 
    if available then
      Look up  $c_{prob}$ 
      Look up search positions room  $r_{sp}$ 
    end if
    if  $r_{user} \neq r_{sp}$  then
       $c_{penalty} = 0$ 
    else
       $c_{penalty} = 1$ 
    end if
    if  $c_{obj} < c_{min}$  then
       $c_{min} = c_{obj}$ 
    end if
  end for
  Move to search position
  Remove location from list of search positions
  Start object recognition
  if object is recognized then
    Set object located to TRUE
    Grasp object
  end if
  if Grasp succeeded then
    Put object on tray
  end if
  Update probabilities of all recognized objects
  Move to the user
  Inform the user about the object location
end while

```

Fig. 3: Algorithm: locate object procedure

B. On the fly extraction of search locations

Our proposed approach automatically generates search positions on the fly after entering a room and executing the algorithm in Fig. 4 which uses semantic segmentation, returning at least 2 search positions or none at all (when no table planes were detected).

```

Move to center pose in a room
while rotation  $< 360^\circ$  do
  Rotate  $30^\circ$  ccw
  Call semantic segmentation service
  Receive possible search poses
  Remove poses from outside of the known map and
  outside the current room
end while

```

Fig. 4: Algorithm: obtain search poses

Semantic segmentation: Our proposed point cloud processing pipeline consists of four steps, depicted in Fig. 5. First, we create an over-segmentation of the scene, clustering it into many small homogeneous patches. In the second step, we compute a manifold but efficient-to-compute feature set for each patch. The resulting feature vector is then processed by a classifier, which yields a probability for each patch being assigned a specific label. To that end, we use a randomized decision forest, a classifier which is intensively discussed in [17]. We train the classifier on the publicly available NYU Depth V2 dataset [18], containing 1,449 indoor frames, recorded by a Microsoft Kinect and densely labelled with more than 1,000 labels. In the last stage of our processing pipeline the classification results set up a pairwise Markov Random Field (MRF), whose optimization yields the final labelling. This last step smoothes the labelling out to correct ambiguous classification results due to noisy local patch information. The final labelling then corresponds to the Maximum-a-posteriori of the output of the MRF. In particular, because our robot can only grasp objects located on tables, we only consider positions next to large clusters of points labelled “table” as suitable positions to detect (and consequently grasp) objects. Therefore, after calculating the semantic labels of the current scene, we use simple Euclidean clustering to obtain all tables in the scene. The resulting search positions are then defined by a simple heuristic, which is explained in Fig. 6. For further details about our semantic segmentation pipeline we refer to [19].

From this method we gain the knowledge of search locations inside the map, which provide a horizontal plane in a certain height range (to filter out the floor and the lower surface of wall-mounted cupboards or the ceiling). As these locations depend on the first and second principal components derived from the segmented planes we have to filter out positions which are not reachable (e.g. inside of another piece of furniture, or outside of the map) as well as locations behind a wall. For this we first check if the detected table surfaces and the robot are within the same room boundaries and remove any search locations which do

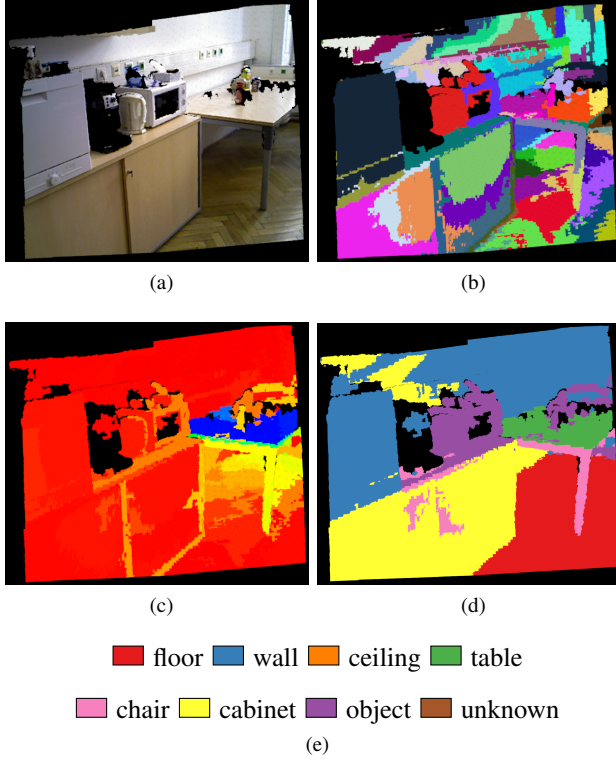


Fig. 5: Intermediate steps of our segmentation pipeline (a) input image (b) oversegmentation (c) conditional label probabilities (here for label *table*, red=0, blue=1) (d) final result after MRF. (e) Colour code of final result




Fig. 6: Heuristic to define search positions for a table cluster projected on the ground-plane (red points). Blue lines: Principal axes of the cluster. The search positions (green dots) are placed on the second principal axis, adhering a security distance d from the edge of the table.

not fulfil this criteria. Secondly we do not consider positions that are inside of an occupied grid cell of the static map.

C. Object recognition

After the robot reaches a “search position” an attempt to recognize the object inside the scene and obtain the objects 6-DOF pose (position and orientation). The recognizer is called with the RGB-D data collected by the Kinect on the robot’s head. It combines three different recognition pipelines (2D and 3D) and merges the results during the verification stage. For details of the recognition process we refer to [20]. Results from successfully recognized objects are the names as well as the poses of all objects. Prerequisite for the recognition process is a trained model of every object which should be localized.

D. Inform the user

After the robot recognizes the object at a certain pose the user has to be informed about the current location, therefore the “search user” procedure is activated. This procedure will plan a path from the current room to the one in which the user has been detected the last time (either by the robot itself or by a motion detection system), move there and attempt to detect the user by combining 2D face detection technique, introduced by Viola and Jones [21] as well as 3D detection and tracking of human body parts [22], [23]. If the user is not detected the search will continue until all rooms have been visited by the robot. After detection the user gets informed about the location of objects, during which we use the annotation of the rooms in a reverse order to give the user a pose in human understandable form instead of the 6-DOF pose values.

V. EXPERIMENTS

All experiments were conducted on the first prototype of our custom made HOBbit platform [24] Fig. 7. Its base is a mobile platform with two drive motors with an integrated gear and attached wheels with differential drive. The perception system consists of one RGB-D camera (ASUS Xtion Pro Live) mounted ground-parallel at a height of 40 cm and one RGB-D camera (Kinect) mounted at 124 cm height as part of the robot’s head on a pan-tilt unit. The lower camera is used for self-localization, the upper one for user and object detection, gesture and object recognition and both cameras contribute to the obstacle detection during navigation. For manipulation PT1 is equipped with an IGUS Robolink arm with a 2-finger Finray-based gripping system. Further parts of the platform include a 15 in tablet computer for touch input and speech recognition as well as a 7 in display for the purpose of displaying the robot’s face. To stash away retrieved objects a tray is mounted behind the tablet and above the robot arm.

To evaluate and compare the two implementations described in Section IV we conducted four tests, each consisting of one execution of both implementations. The object that had to be located was placed according to table I. As objects we used an Asus Xtion Pro box, a brown handbag and



Fig. 7: HOBBIT platform - prototype 1

an orange wallet, which were placed at the same locations for the corresponding search runs. The robot was placed at the same start location before each pair of runs, with a total of 4 different start locations. For each run we collected the following data:

- The duration of the entire run until the robot detected the object or ended the search run.
- The number of “search positions”. (Pre-defined, calculated and with outliers removed)
- The success rate of the search run. (Was the object recognized on the actual location?)

| object | start | location | object detected? | | duration [min]:[sec] | | # |
|----------|--------------|--------------------|------------------|----|----------------------|-------|----|
| | | | P | S | P | S | |
| Asus box | Kitchen | Dining room table | Y | Y | 7:12 | 16:31 | 9 |
| Handbag | Living room | Bedroom nightstand | N | N* | 9:52 | 12:15 | 17 |
| Wallet | Central hall | Living room shelf | Y | N | 7:04 | 23:07 | 36 |
| Asus box | Dining room | Kitchen cupboard | N | Y | 8:51 | 10:45 | 17 |

TABLE I: Test run set-up and results. #...number of “search positions”, N*...false positive recognition of the object, P...pre-defined locations, S...semantic segmentation-based

VI. DISCUSSION AND OUTLOOK

At a first sight the results of our approach uses up to 2.5 times the time of the pre-programmed search task, which was expected due to the fact that the robot not only searches on pre-defined tables but also on window sills, shelves, etcetera. We further identified 3 major time-costing culprits in our system and set-up. First is the sequential order of the semantic segmentation and the 30°-rotation plus the point

cloud acquisition after the entrance of a room. As they both take almost the same time to complete (5 sec and 4 sec) we can save 48 sec for each room. In our 5 virtual room set-up this would lead to a reduction of up to 4 min, depending on the number of actually visited rooms. This leaves us room for future improvements for our proposed approach. The third culprit is the higher number of “search positions” (up to 20) that are inside of the desired room opposed to only one, which leads to a higher probability to find an object at the cost of increased search duration. To reduce this number of positions it is possible to cluster these poses together under the premise that they are close enough and the field of view will overlap at least a certain amount so that a possibly located table plane is not discarded. Fig. 8 shows the result of search position calculations without the virtual room constraints and a number of poses which could be clustered together, especially on the left side and in the right handed lower corner.

Another approach is to combine the semantic search algorithm with the pre-defined locations where the “search locations” are autonomously generated once instead of the manual labelling. This would improve the speed of the whole search procedure but remove any flexibility towards changes in the environment.




Fig. 8: Manual overlay of 12 search position calculations in our lab

VII. CONCLUSION

In this paper we presented a SMACH-based behaviour control for more intelligent indoor fetch-and-carry tasks by the means of semantic segmentation. Our approach does not require manually annotated search positions and thereby increases flexibility and human-oriented perception. We could show the feasibility of the approach by the means of a first series of experiments and outlined suggestions how it can be further improved. After implementing these improvements, we will perform a user study to assess how our approach is evaluated in terms of perceived intelligence and overall satisfaction by naive users.

ACKNOWLEDGMENT

This work has been partially funded by the European Commission under FP7-IST-288146 HOBbit.

REFERENCES

- [1] W. Burgard, A. B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun, "The interactive museum tour-guide robot," in *AAAI/IAAI*, 1998, pp. 11–18.
- [2] S. Thrun, M. Bennewitz, W. Burgard, A. B. Cremers, F. Dellaert, D. Fox, D. Hahnel, C. Rosenberg, N. Roy, J. Schulte *et al.*, "Minerva: A second-generation museum tour-guide robot," in *In Proceedings of IEEE International Conference on Robotics and Automation (ICRA) 1999*, vol. 3. IEEE, 1999.
- [3] H.-M. Gross, H.-J. Böhme, C. Schröter, S. Mueller, A. König, C. Martin, M. Merten, and A. Bley, "Shopbot: Progress in developing an interactive mobile shopping assistant for everyday use," in *IEEE International Conference on Systems, Man and Cybernetics (SMC) 2008*. IEEE, 2008, pp. 3471–3478.
- [4] B. Graf, U. Reiser, M. Hagele, K. Mauz, and P. Klein, "Robotic home assistant care-o-bot® 3-product vision and innovation platform," in *IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO), 2009*. IEEE, 2009, pp. 139–144.
- [5] M. Lohse, "The role of expectations in hri," in *New Frontiers in Human-Robot Interaction*. SSAISB (ISBN 190295680X), 2009.
- [6] RoboCup@Home Technical Committee *et al.*, "Rules & regulations. final version," 2009.
- [7] J. Bohren, R. B. Rusu, E. G. Jones, E. Marder-Eppstein, C. Pantofaru, M. Wise, L. Mosenlechner, W. Meeussen, and S. Holzer, "Towards autonomous robotic butlers: Lessons learned with the pr2," in *IEEE International Conference on Robotics and Automation (ICRA) 2011*. IEEE, 2011, pp. 5568–5575.
- [8] J. Stuckler and S. Behnke, "Benchmarking mobile manipulation in everyday environments," in *IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO), 2012*. IEEE, 2012, pp. 1–6.
- [9] D. Meger, P.-E. Forssén, K. Lai, S. Helmer, S. McCann, T. Southey, M. Baumann, J. J. Little, and D. G. Lowe, "Curious george: An attentive semantic robot," *Robotics and Autonomous Systems*, vol. 56, no. 6, pp. 503–511, 2008.
- [10] H. Huttenrauch and K. Severinson Eklundh, "Fetch-and-carry with cero: observations from a long-term user study with a service robot," in *11th IEEE International Workshop on Robot and Human Interactive Communication*, 2002. IEEE, 2002, pp. 158–163.
- [11] M. L. Walters, K. Dautenhahn, S. N. Woods, and K. L. Koay, "Robotic etiquette: results from user studies involving a fetch and carry task," in *2nd ACM/IEEE International Conference on Human-Robot Interaction (HRI) 2007*. IEEE, 2007, pp. 317–324.
- [12] T. Baier, M. Huser, D. Westhoff, and J. Zhang, "A flexible software architecture for multi-modal service robots," in *Multiconference on Computational Engineering in Systems Applications, IMACS*, vol. 1. IEEE, 2006, pp. 587–592.
- [13] E. Pot, J. Monceaux, R. Gelin, and B. Maisonnier, "Choregraphe: a graphical tool for humanoid robot programming," in *The 18th IEEE International Symposium on Robot and Human Interactive Communication, (RO-MAN) 2009*. IEEE, 2009, pp. 46–51.
- [14] S. Kopp, B. Krenn, S. Marsella, A. N. Marshall, C. Pelachaud, H. Pirker, K. R. Thórisson, and H. Vilhjálmsson, "Towards a common framework for multimodal generation: The behavior markup language," in *Intelligent virtual agents*. Springer, 2006, pp. 205–217.
- [15] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "Ros: an open-source robot operating system," in *ICRA workshop on open source software*, vol. 3, no. 3.2, 2009.
- [16] F. Siepmann, "Behavior coordination for reusable system design in interactive robotics," 2013.
- [17] A. Criminisi and J. Shotton, Eds., *Decision Forests for Computer Vision and Medical Image Analysis*. Springer London, 2013. [Online]. Available: <http://link.springer.com/10.1007/978-1-4471-4929-3>
- [18] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus, "Indoor segmentation and support inference from RGBD images," in *ECCV*, 2012.
- [19] D. Wolf, M. Bajones, J. Prankl, and M. Vincze, "Find my mug: Efficient object search with a mobile robot using semantic segmentation," in *38th Annual Workshop of the Austrian Association for Pattern Recognition (ÖAGM)*, 2014.
- [20] A. Aldoma, F. Tombari, J. Prankl, A. Richtsfeld, L. Di Stefano, and M. Vincze, "Multimodal cue integration through hypotheses verification for rgb-d object recognition and 6dof pose estimation," in *IEEE International Conference on Robotics and Automation (ICRA) 2013*. IEEE, 2013, pp. 2104–2111.
- [21] P. Viola and M. Jones, "Robust real-time face detection," *International Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154, 2004.
- [22] N. Grammalidis, G. Goussis, G. Troufakos, and M. G. Strintzis, "3-d human body tracking from depth images using analysis by synthesis," in *International Conference on Image Processing, 2001*, vol. 2. IEEE, 2001, pp. 185–188.
- [23] J. Shotton, T. Sharp, A. Kipman, A. Fitzgibbon, M. Finocchio, A. Blake, M. Cook, and R. Moore, "Real-time human pose recognition in parts from single depth images," *Communications of the ACM*, vol. 56, no. 1, pp. 116–124, 2013.
- [24] D. Fischinger, P. Einramhof, W. Wohlking, K. Papoutsakis, P. Mayer, P. Panek, T. Koertner, S. Hofmann, A. Argyros, M. Vincze, A. Weiss, and C. Gisinger, "Hobbit - the mutual care robot," *ASROB-2013 in conjunction with IEEE/RSJ Intern. Conference on Intelligent Robots and Systems (IROS), Japan (2013)*, 2013.